

# Is HPC Going Green?

*Looking at how we can change the power equation*

► Horst D. Simon, Ph.D.

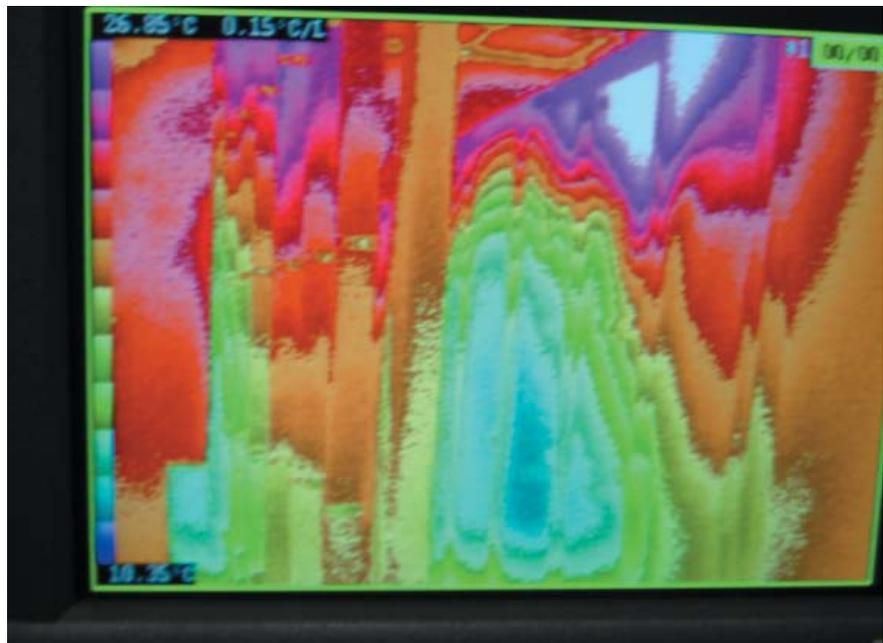
For the past 20 years, even as high performance computing centers have increased significantly in number, size and power requirements, the HPC community has until recently held the attitude: Why do we need to be energy efficient? We can just crank in more megawatts.

But in the United States and around the world, we're changing our thinking. Call it "Green Computing" or "Energy Efficient Computing," the problem facing all of us is that, with the rising cost of electricity and with increasing demand for both powering and cooling our computing systems, power costs now loom larger than the cost of hardware for supercomputing centers. A study commissioned by the U.S. Environmental Protection Agency estimates that power consumption by servers doubled between 2000 and 2005 worldwide, and that the total amount of electricity consumed by servers worldwide now costs about US \$7.2 billion. This is nearly the same order of magnitude as today's investment in HPC technology (US \$9.2 billion). In short, we have reached a critical threshold that should give us cause to consider the question of power consumption as a potentially limiting factor to future growth in HPC.

While the energy problem facing centers in the U.S. is daunting, during a visit by some of our European colleagues last year, we learned that they are paying nearly 10 times as much for power. In fact, a recent survey by IDC found that an overwhelming majority of facilities managers named power and cooling as the most pressing issues of concern to them. A study of exaflops computing in the



**FIGURE 1:** Using a thermal imaging camera to acquire very detailed measurements of the temperature of the airflow in an HPC machine room. The wooden frame can be moved to any part of the aisle to measure the air temperature (reflected back to them by the white strips) from any part of the aisle.



**FIGURE 2:** The thermal imaging camera provides far higher resolution than could be accomplished using rack-embedded sensor meshes or spot thermal meters.

## High Performance Computing

### The International Supercomputing Conference

The International Supercomputing Conference (ISC) began in 1986 as an event offering both vendors and researchers an international perspective on high performance computing. The series began in 1986, when about 80 participants gathered in Mannheim, Germany. It then moved to Heidelberg in 2001 and grew to become Europe's leading supercomputing conference, providing a unique atmosphere for collaboration and cooperation and connecting the HPC community on a global scale.

Founded by Prof. Dr. Hans Werner Meuer, who has served as conference chair since the beginning, the conference relocated to historic Dresden in 2006, where it was able to accommodate more attendees and exhibitors in the state-of-art International Congress Center ([www.dresden-congresscenter.de/eng/eng\\_1.htm](http://www.dresden-congresscenter.de/eng/eng_1.htm)). The move to Dresden proved to be a good one: attendance jumped to 915 and 1,213 in the years 2006 and 2007 respectively, and the number of international exhibitors grew from 48 in 2005 (Heidelberg) to 74 in 2006 and 85 in 2007 (both Dresden).

In 2008, ISC will continue to offer a strong three-day technical program, featuring experts from research and industry, paralleled by a huge exhibition with the key players in the HPC market. Special events will be aimed at both industry participants (Industrial Session, Hot Seat Session and Exhibitor Forum) and research scientists (Scientific Day, Research Posters, Birds-of-a-Feather (BoF) sessions).

For more information: [www.supercomp.de/isc08](http://www.supercomp.de/isc08)

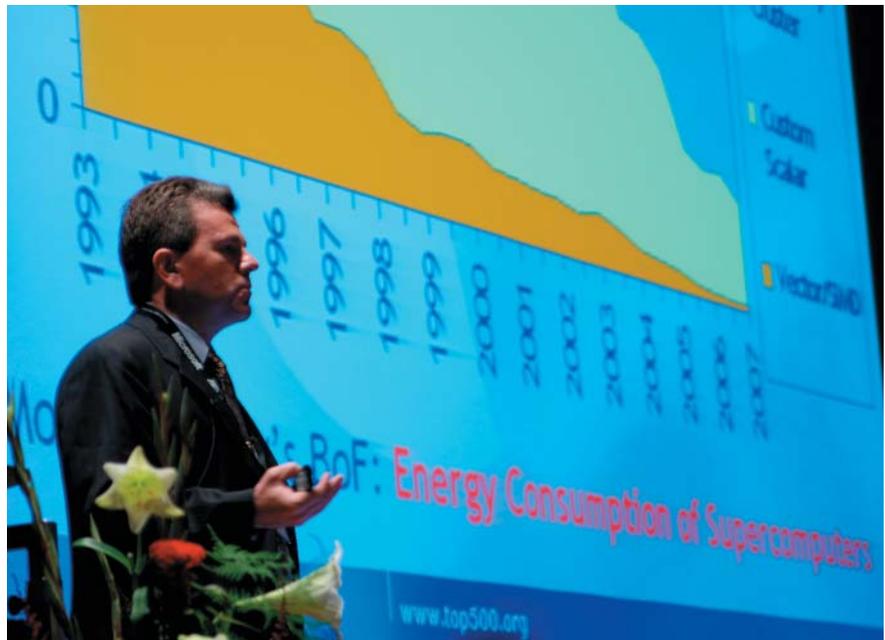


U.S. came to the conclusion that, by projecting today's technology, an exaflops computer might require 120 MW of power.

This is clearly a problem which must be addressed by the HPC community as a whole — trying to solve it center by center will be an inefficient use of our resources. One key step in the sharing of information and expertise will be a panel discussion at the 2008 International Supercomputing Conference being held June 17-20 in Dresden, Germany. As a truly global meeting of the HPC community, ISC is an ideal forum for exchanging ideas on this important topic.

In the discussion of “The Greening of HPC — Will Power Consumption Become the Limiting Factor for Future Growth of HPC?,” an international panel of experts will address the questions of what are the power limitations of current technology, and how can we change the equation to assure the future rapid growth of HPC performance without contributing even more to carbon emissions and global warming?

I am pleased to have been asked by ISC General Chair Professor Hans Meuer to chair this session and have enlisted a panel with diverse expertise in tackling this problem at various



**FIGURE 3:** Eric Strohmaier speaks about energy usage during his TOP500 presentation at ISC'07. Courtesy of Tim Krieger/ISC

levels. Such approaches range from studying air temperature at different locations in a machine room, to better understanding opportunities for more efficient cooling, to using computational fluid dynamics to model cooling flows, to designing entirely new architectures with the goal of reducing power consumption from the very outset.

Another session at ISC '08 will fea-

ture a significant contribution to more fully understanding power consumption and measurement for HPC systems. For the first time, the TOP500 list also will include information about electrical usage and assess this in terms of system performance. Vendors are being asked to provide data on power consumption while running Linpack, and we are looking forward to their submissions. This information

will be combined with data on overall performance and memory to generate a “utility” metric.

Erich Strohmaier, one of the founding editors of the TOP500 list, says this utility metric will be used to answer questions such as “When is my new computer system twice as useful as my old one?” If the old system is just doubled in terms of performance, memory and power consumption, is it really twice as good as the old version? However, if you can get twice the performance

and memory, but consume the same amount of power, then is the advantage clear? This and other questions will be discussed during the TOP500 presentation in the opening session.

It appears that the end of Moore’s Law might very well be determined by our ability to cool our computer systems and not by our ability to shrink feature sizes on chips. The total cost of power consumption can be much higher than electric charges alone, especially when you consider the expense of enlarging the existing

infrastructure to accommodate systems well beyond power and size envelopes of current systems. Hopefully, the sessions at ISC ’08 will add significantly to the continuing discussion of this critical issue facing the global HPC community. **SC**

*Horst Simon is Associate Laboratory Director for Computing Sciences at Lawrence Berkeley National Laboratory and is one of four editors of the “TOP500” list of the world’s most powerful computing systems.*

### The Greening of HPC – International Panel of Experts

► **Wu-chun Feng** is an associate professor of Computer Science and Electrical and Computer Engineering and also directs the Synergy Laboratory at Virginia Tech in the U.S. Wu has been tackling this issue since before it became popular, having introduced the controversial notion of “low power” to the world of supercomputing at the SC’01 conference. Shortly thereafter, he debuted Green Destiny, a 240-node cluster supercomputer in five square feet that consumed a mere 3.2 kilowatts of power (when booted diskless). He is also the author of more than 100 articles in peer-reviewed technical publications in high-performance networking and computing, high-speed systems monitoring and measurement, low-power and power-aware computing, and bioinformatics.

► **John Gustafson** has 33 years of experience using and designing compute-intensive systems, including the first matrix algebra accelerator and the first commercial massively-parallel cluster while at Floating Point Systems. His pioneering work on a 1024-processor nCUBE at Sandia National Laboratories created a watershed in parallel computing, for which he received the inaugural Gordon Bell Award. He also has received three R&D 100 Awards for innovative performance models, including the model commonly known as Gustafson’s Law or Scaled Speedup. Most recently, Gustafson has led high-performance computing efforts at ClearSpeed and Sun Microsystems.

► **Satoshi Matsuoka** is a full professor at the Global Scientific Information and Computing Center (GSIC) of Tokyo Institute of Technology, where he leads the Research Infrastructure Division Solving Environment Group. He pioneered grid computing research in Japan in the mid-90s along with his collaborators, and currently serves as sub-leader of the Japanese National Research Grid Initiative (NAREGI) project, which aims to create middleware for next-generation CyberScience Infrastructure. He was also the technical leader in the construction of the TSUBAME supercomputer, which has become the fastest supercomputer in Asia-Pacific, performing at 85 teraflops (peak) and placing 38.18 Teraflops (Linpack, 7th on the June 2006 list).

► **Franz-Josef Pfreundt** brings a diverse set of research interests to the discussion, including fluid dynamics, porous media, image analysis and parallel computing. He is the division director and head of the “Competence Center for HPC and Visualization” at the Fraunhofer Institute for Industrial Mathematics (ITWM) in Kaiserslautern, Germany. In 2001, he was awarded the prestigious Fraunhofer Research Prize, along with Konrad Steiner and their research group for their work on microstructure simulation.

► **John Shalf** is the leader for the Science Driven Systems Architecture group at the National Energy Research Supercomputing Center (NERSC) at Lawrence Berkeley National Laboratory. Shalf is a member of a team investigating the use of low-power, high-performance multicore chips to design a prototype computer architecture for studying climate change. His interests include computer architecture, performance evaluation/benchmarking, parallel I/O, high performance networking and data analysis frameworks.

### New Products

#### SiCortex SC5832

SiCortex SC5832 is part of the SiCortex family of ultra low power high performance Linux systems. It is a complete 5.8 TFlops computing system packaged in a single cabinet, which holds 36 circuit modules. Each circuit module accommodates 27 node chips and associated memory. There are two memory sockets per node, which can be configured with up to 8 GByte per node, for a system total of 8 TBytes.



SiCortex, [www.sicortex.com](http://www.sicortex.com)

#### Relion 1670

The Relion 1670 server is based on the Intel 5400 ‘Seaburg’ chipset platform and Intel Xeon 5200 and 5400 series processors. The server is rack-mountable and supports a higher 1600 MHz front side bus speed whereas typical mainstream enterprise servers are limited to 1333 MHz. The system also provides up to 128 GB of main memory, a unique capability in a compact 1U form factor. It provides management and serviceability features to support large HPC clusters and mission critical enterprise applications, and supports free and commercial versions of Linux.



Penguin Computing,  
[www.penguincomputing.com](http://www.penguincomputing.com)